

Analysis of KEK-ATF Optics and Coupling Using LOCO

A. Wolski

Lawrence Berkeley National Laboratory

J. Nelson, M. Ross, M. Woodley

Stanford Linear Accelerator Center

S. Mishra

Fermi National Accelerator Laboratory

January 12th, 2004

Abstract

LOCO is a code for analysis of the linear optics in a storage ring based on the closed orbit response to steering magnets. The analysis provides information on focusing errors, BPM gain and rotation errors, and local coupling. Here, we report the results of an application of LOCO to the KEK-ATF. Although the analysis appears to have provided useful information on the optics of the machine, it appears that one of the main aims of the study – to reduce the vertical emittance by correcting the local coupling – was not successful, and we discuss some possible reasons for this.

1 Introduction

The ATF has recently reported a vertical emittance of below 5 pm [1]. Although this meets the specification for the damping rings for an X-band linear collider, detailed studies of potentially limiting effects in beam dynamics, and the development of high performance diagnostics, motivate efforts to reduce the vertical emittance to even smaller values. This would require very precise correction of betatron coupling and vertical dispersion. Simulation studies have shown that it may be possible to achieve a vertical emittance close to 1 pm [2]. The Advanced Light Source (ALS) at LBNL has also achieved a vertical emittance of around 5 pm [3]. In this case, the coupling correction was achieved in a reasonably straightforward manner using the optics analysis code LOCO [4] to determine the appropriate strengths of the skew quadrupoles for correcting the local coupling.

In this note, we report the results of an attempt to apply LOCO to the ATF. The aims of the study were:

- to obtain information on the optics of the ATF, in particular to identify any focusing errors arising from variation of quadrupole strengths from nominal values;
- to obtain information on the diagnostics system of the ATF, in particular to estimate BPM gains and coupling errors;
- to reduce the vertical emittance of the ATF by determining appropriate settings for the skew quadrupoles in a manner similar to that used at the ALS.

In using LOCO, one provides as input measurements of the response matrix between the steering magnets and the BPMs, and measurements of the horizontal and vertical dispersion. LOCO then adjusts parameters in a lattice model to reproduce the input data. To the extent that the response matrix of the fitted model matched the measured response matrix, it appears that the first two of the above aims were achieved. However, estimates of the change in vertical emittance based on beam lifetime measurements suggested that our attempt at correcting the coupling actually *increased* the vertical emittance. In this note, we present and discuss the results of the LOCO analysis at the ATF, and consider the possible reason for our failure to correct the coupling.

2 Measurement Procedure

The ATF contains 47 horizontal orbit correctors, and 50 vertical orbit correctors. There are 96 BPMs in each plane. Data collection was semi-automatic, with each orbit corrector varied in steps, with a number of orbit measurements made at each step. The usual operating mode of the ATF is to inject a fresh bunch train at a rate of around 1.6 Hz, and record the orbit on a specified turn (normally 500,000 turns after injection). The BPMs are known to have a dependence (varying with time) on the beam current, and orbits were therefore only recorded for current within a specified range. For a small number of correctors, changes in the kick angle distorted the orbit in such a way that injection efficiency was very poor, and these correctors were therefore excluded from the analysis. The raw data consist of a set of BPM readings for each setting of each corrector. The data are preprocessed to give the orbit variation at each BPM with respect to the variation of each corrector – i.e. the closed orbit response matrix. This is the principal input for LOCO.

We note that the BPM resolution is typically a few microns. By this, we mean that repeated orbit measurements over successive injections, with nominally no change in the machine settings, will vary with an rms of a few microns. Thus, averaging over a number of orbits provides data with good resolution for the LOCO analysis. The resolution of each BPM is estimated from the collected data, and is used by LOCO in weighting the fit to the data provided by each BPM.

The horizontal and vertical dispersion were also provided as input data for the LOCO analysis. This breaks a degeneracy that otherwise occurs between the corrector strengths and the BPM gains. We discuss the significance of the vertical dispersion in the analysis below.

Data were collected on two shifts, on December 16 and December 19. Attempts were made to reduce the coupling by adjusting the skew quadrupole strengths on both shifts, but data for estimating the change in vertical emittance (beam lifetime measurements) were collected only on the December 19 shift.

3 Analysis of Focusing Errors, BPM Gains and Corrector Strengths

Here, we present the results of the analysis of data collected on December 19. The data from December 16 give very similar results. The parameters varied in fitting the model to the measured response matrix were as follows:

- BPM gains and couplings;

- corrector magnet kicks and couplings;
- strengths of all quadrupoles;
- strengths of skew quadrupoles superposed on SF sextupoles.

Note that all vertical focusing in the arcs is provided by the gradients in the bending magnets. A set of quadrupoles (QF1R) close to the bending magnets is available to provide focusing adjustments, and these quadrupoles are nominally powered as weak horizontally focusing magnets. We found in the LOCO analysis that the gradients in these quadrupoles are nearly degenerate with the gradients in the dipoles; variations in the focusing of the QF1R quadrupoles therefore include variations in the gradients in the dipoles.

In the ATF, every sextupole has an independently adjustable skew quadrupole trim winding for compensating the coupling. In principle, the skew quadrupole components on the SF and SD sextupoles can be independently fitted as parameters in the LOCO analysis; however the horizontal and vertical phase advances between adjacent SF and SD sextupoles are very nearly equal, resulting in a near degeneracy of the skew trims on these magnets from point of view of the coupling. Attempting independent fits of these parameters resulted in very large values for the skew gradients being found by LOCO, but with adjacent skew quadrupoles nearly canceling out. We therefore used only the skew quadrupoles in one family of sextupoles as parameters in the fit.




Figure 1

Terms in the measured response matrix (left) and difference between the measured response matrix and the modeled response matrix after fitting with LOCO (right).

Figure 1 shows the measured response matrix, and the difference between the measured response matrix and the modeled response matrix, after the model is fitted to the response matrix using LOCO. Note the different scales on the two plots, and that the cross-plane sectors of the response matrix are very much smaller than the in-plane sectors (horizontal BPMs are

numbered 1-96, vertical BPMs 97-192; horizontal corrector magnets are numbered 1-47, vertical corrector magnets 48-97).

The cross-plane sectors are shown separately in Figure 2 and Figure 3. Note the scales on the difference plots are the same as on the plots of the measurements; the residuals of the coupling components after the fit are significant, suggesting that using LOCO to determine a coupling correction could have only limited effect.




Figure 2

Sector of the response matrix corresponding to horizontal BPMs and vertical corrector magnets.




Figure 3

Sector of the response matrix corresponding to vertical BPMs and horizontal corrector magnets.




Figure 4

Distribution of the residuals of the fit to the response matrix, in units of the BPM resolution.

A more detailed indication of the quality of the fit to the response matrix is given in Figure 4. This shows a distribution of the residuals in units of the “resolution” of the associated BPM. In other words, to construct the distribution, we took the difference between the measured orbit response matrix and the orbit response matrix of the fitted LOCO model, divided each value by the resolution of the associated BPM, and plotted a histogram over all elements in the response matrix (normalized by the total number of points).




Figure 5

Nominal and fitted quadrupole gradients. The points indicated with a cross and joined with a red line are the nominal gradients; the blue points with the error bars are the fitted gradients.

The fitted quadrupole gradients are shown in Figure 5; in most cases, the gradient errors appear small, though there is some variation in the QD family of quadrupoles in the arcs (nominally close to zero strength). As we mentioned above, this is probably an indication of variation of the gradients in the dipoles.

The fitted BPM gains are shown in Figure 6. We note that there seems to be a systematic gain error in both the horizontal and vertical BPMs. This may be real, or an artifact from an error in the dispersion data provided. The fitted BPM couplings are shown in Figure 7. The coupling is defined as the *measured* beam motion in one plane resulting from a *real* unit beam motion in the other plane. If the coupling of a particular BPM arises from rotation of that BPM around the beam axis, then a coupling of 0.01 would correspond to a rotation of 10 mrad. Of course, there is no guarantee that the source of coupling is BPM rotation; in LOCO, the horizontal and vertical couplings are independent parameters.




Figure 6
BPM gains fitted by LOCO.




Figure 7
BPM couplings fitted by LOCO.




Figure 8

Nominal and fitted corrector kicks. The points indicated with a cross and joined with a red line are the nominal kicks; the blue points with the error bars are the fitted kicks.

Figure 8 shows the fitted corrector kicks compared to the nominal values. Note that some corrector magnets were excluded from the data because changes in their strengths resulted in poor injection efficiency. We see that there appears to be a systematic error in the kick strengths, with the fitted kicks generally *larger* than the nominal values. For the BPM gains, we observed a systematic error, with the fitted gain generally *less than* one. This is suggestive of some error in the dispersion data provided, although it is not clear how this could come about.




Figure 9

Horizontal and vertical “dispersion” (actually change in closed orbit with respect to change in RF frequency). The measured values are shown as crosses in red; the fitted values from LOCO are shown as circles in blue.

The measured and fitted horizontal and vertical dispersion are shown in Figure 9. In the horizontal dispersion, there is a large discrepancy between the measured and fitted dispersion near the middle of one of the arcs. This is consistent between different sets of data, and the reason for it has not been identified.

The fit to the vertical dispersion was carried out with a weight of unity (relative to the horizontal dispersion and response matrix fits). This is probably not the correct weight to use. Although it appears a reasonably good fit is obtained to the vertical dispersion, it assumes that all the vertical dispersion is generated by coupling in the skew quadrupoles. LOCO does not address the effects of closed orbit distortions (except indirectly through the effect of feed-down focusing terms resulting from beam offsets in the sextupoles). If vertical steering does make a significant contribution to the vertical dispersion, then fitting the vertical dispersion with a large weight will distort the values found for the skew quadrupole strengths, compromising the fit to the coupling parts of the response matrix.

4 Analysis of Skew Quadrupole Strengths and Coupling Correction

4.1 Analysis of Data from December 16

During the shift on December 16, we attempted to apply a coupling correction iteratively. The expected dominant sources of coupling in the machine are rotations of the quadrupoles about the beam axis, and vertical offset of the beam in the sextupoles. To compensate these sources of coupling, the sextupoles are fitted with skew quadrupole trim windings, which can all be independently adjusted. Coupling correction involves determining the optimum settings for the skew quadrupoles to compensate the quadrupole rotations and sextupole offsets. The procedure we adopted was as follows:

1. Collect data for LOCO, and carry out an analysis assuming that all coupling comes from skew quadrupoles superposed on the SF sextupoles.
2. Adjust the strengths of the skew quadrupole trim windings on the SF and SD sextupoles to cancel the skew quadrupole strengths found from step 1.
3. Repeat the above steps as time allows, or no further changes to the skew quadrupole strengths are needed.




Figure 10

Change in fitted skew quadrupole strengths as determined by LOCO analysis of response matrix data, compared with known changes in current in the skew quadrupole trim windings on the sextupoles.

A key indication of the ability of LOCO to correct the coupling is the change in the skew quadrupole strengths found by LOCO between two iterations, compared with the known changes applied to the currents in the skew quadrupole trim windings on the sextupoles. Such a comparison is shown in Figure 10. The data are taken from the machine in its initial condition and after the first applied correction.

We observe that the fitted changes in skew quadrupole strengths are well correlated with the known changes in the current in the trim windings. This gives us confidence that LOCO is providing meaningful results.

The shift on December 16 allowed time for three sets of LOCO data to be collected, and two attempts at coupling correction to be made (based on the first and second set of data). The distribution of components in the coupling sectors of the response matrix on each iteration are shown in Figure 11.




Figure 11

Distribution of coupling components in the response matrix, before and after attempts at coupling correction.

We observe that after the first skew correction, there is some reduction in the width of the distribution, with a significant increase in the peak. This implies that there has been a general reduction in the size of the coupling components in the response matrix; the increase in the peak of the distribution may be explained if many points in the tails of the distribution have been significantly reduced.

After the second skew correction, it appears that there has been an *increase* in the coupling as indicated by the response matrix.




Figure 12

Skew quadrupole strengths as found by LOCO before skew correction, and after the first and second attempts at skew correction.

The skew quadrupole strengths found by LOCO before and after skew correction are shown in Figure 12. Note that the current in one of the trim windings (SF1.10) was at its limit, with no further correction possible. It appears that the first correction was successful in reducing the strengths of the skew quadrupoles. The second iteration did not appear to lead to any further reduction in the skew quadrupole strengths.

Figure 13 shows the measured vertical dispersion before and after skew correction. The skew correction appears to *increase* the vertical dispersion. It is possible that this is consistent with a

reduction in the skew coupling, if the orbit has large vertical steering, in which case the skew coupling might act to compensate the dispersion generated by steering; we admit that this is not a likely scenario.




Figure 13

Measured vertical dispersion before and after skew correction.

The results from the response matrix measurement, and the fit to the skew quadrupole strengths are consistent in indicating a reduction in the coupling following the first skew correction, and no further reduction (and possibly some increase) after the second skew correction. The reasons for the failure of the second skew correction are unclear. Figure 10 suggests that the skew quadrupole strengths being found are meaningful, but the scatter in this correlation plot indicates that a resolution limit is being reached after the first correction (the residual strengths of the skew quadrupoles fitted by LOCO are comparable to the deviations of the points from the fitted line in the correlation plot). It is also possible that the vertical dispersion was being fitted with too large a weight with the result that the effects of vertical steering “contaminated” the fitted strengths of the skew quadrupoles.

4.2 Analysis of Data from December 19

Because of time limitations, only a single attempt at coupling correction was possible. A response matrix measurement was made and the data analyzed with LOCO to determine the

required changes to the currents in the skew quadrupole trim windings. We had hoped to make a direct measurement of the change in the vertical emittance using the laser wire; unfortunately, a hardware failure meant that this was not possible. Instead, we measured the beam lifetime to indicate any change in the vertical emittance.




Figure 14

Skew quadrupole strengths found from LOCO analysis of nominally identical lattices on December 16 and December 19.

Figure 14 compares the skew quadrupole strengths found from the LOCO analysis of the data collected on December 16 and December 19, both before any skew correction. The configuration of the December 16 lattice was reloaded before the December 19 measurements, so the lattices were nominally identical. Mostly, the changes in the fitted skew quadrupole strengths are small, except in one arc.

Figure 15 compares the measured vertical dispersion on December 16 and December 19, in each case before any skew correction was applied. Again, the agreement is reasonably good.




Figure 15

Measured vertical dispersion on December 16 and December 19, before skew correction.




Figure 16

Measured vertical dispersion on December 16 and December 19, after skew correction.




Figure 17

Change in measured vertical dispersion on December 16 and December 19, as a result of skew correction.

Unfortunately, because of time limitations, it was not possible to measure a response matrix for LOCO analysis after the skew correction. However, we did collect dispersion and beam lifetime data. Figure 16 shows the measured vertical dispersion after the skew correction, on December 16 and December 19. Since the skew quadrupole strengths on each date were found to be about the same, approximately the same skew correction was applied on each date. However, there does appear to be some difference between the two dates, in the measured change in vertical dispersion following the skew correction (Figure 17). The reasons for this, and its significance, are unclear.

The results of the beam lifetime measurements are shown in Figure 18. For these measurements, a single bunch was stored in the ring, and the current measured as a function of time. Applying the coupling correction determined from the LOCO analysis clearly *increases* the beam lifetime. Assuming that the Touschek effect is a significant contribution to the beam lifetime, this strongly suggests that the coupling correction actually *increased* the vertical emittance.




Figure 18

Decay of beam current with time for different skew quadrupole settings. The steepest decay (shortest lifetime) is with the skew quadrupoles at their nominal settings before coupling correction; the slowest decay (longest lifetime) is with the skew quadrupoles turned off.

5 Comments and Conclusions

We feel that application of LOCO to the ATF was able to provide some information about the lattice and the diagnostics. The results shown in Figure 10 do indicate that the analysis is able to identify real changes to magnet strengths. Nevertheless, our attempts to reduce the vertical emittance by applying a coupling correction with skew quadrupole strengths determined from the LOCO analysis met with no success. The reason for this is possibly our inattention to the weight applied to the vertical dispersion in fitting the model to the measured data. Steier observed in studies of the ALS [5], that in correcting the coupling, it is important to set the weight for the vertical dispersion correctly: if the weight is too low, then the vertical dispersion remains large after the correction and dominates the emittance; if the weight is too high, then the fitted skew quadrupole strengths are incorrect, because of the effects of vertical steering. Unfortunately, we did not have time in advance of the experimental studies for a proper simulation study to determine the best weight for fitting the vertical dispersion. We should therefore suggest that such studies be carried out, before making any further attempts to apply this correction strategy to the ATF.

We would also comment that the vertical emittance of the ATF is already remarkably small. Depending on the exact errors present in the machine, it is possible that the optimum skew quadrupole settings for low vertical emittance, providing a balance between vertical dispersion and betatron coupling, have already been found. Changes to the skew quadrupole strengths would then result in reducing the betatron coupling at the expense of increasing the vertical dispersion, or vice-versa. In this case, it would be important to reduce the vertical dispersion generated by vertical steering, for example by using the results of beam-based alignment, before applying LOCO to determine changes in the skew quadrupole strengths needed to correct the coupling.

6 Acknowledgements

We thank the ATF group at KEK for their invaluable support, and for their hospitality during our visit. Thanks also to Christoph Steier and James Safranek for support and advice with using LOCO.

7 References

- [1] Y. Honda et al, “Achievement of Ultra-Low Emittance Beam in the ATF Damping Ring”, to be published in Physical Review Letters, 2004.
- [2] K. Kubo, “Simulation Study for Smaller Emittance in ATF Damping Ring”, presented at ISG XI, KEK, 2003. <http://www-proc.kek.jp/ISG11.htm>
- [3] C. Steier et al, “Coupling Correction and Beam Dynamics at Ultralow Vertical Emittance in the ALS”, Proceedings PAC 2003.
- [4] LOCO is a code developed by J. Safranek, and adapted to run under MATLAB by G. Portmann. <http://ssrl.slac.stanford.edu/loco/>
- [5] C. Steier, private communication, 2003.